

Douglas W. Gage

RDT&E Division, Naval Command Control and Ocean Surveillance Center
NCCOSC RDTE DIV 531, San Diego, CA 92152-7383

ABSTRACT

In two previous papers we explored some of the systems aspects of applying large numbers of inexpensive robots to real world applications. The concept of *coverage* can help the user of such a system visualize its overall function and performance in mission-relevant terms, and thereby support necessary system command control functions. An important class of coverage applications are those that involve a *search*, in which a number of searching elements move about within a prescribed search area in order to find one or more target objects, which may be stationary or mobile. A simple analytical framework was employed in the previous work to demonstrate that the design of a cost-effective many-robot search system can depend sensitively on the interplay of sensor cost and performance levels with mission-specific functional and performance requirements. In the current paper we extend these results: we consider additional measures of effectiveness for area search systems to provide a broader basis for a tradeoff of coordinated versus random search models, and we explore how to deliberately achieve effectively randomized search strategies that provide uniform search coverage over a specified area.

1. INTRODUCTION

The rapid evolution of micromechanical fabrication techniques and other sensor, effector, and control/processing technologies will soon make possible the development of very inexpensive autonomous mobile robots with fairly limited sensor capabilities [1]. While the popular and semi-popular press has discovered the opportunity of addressing real world applications with "swarms" of "insect robots" [2,3,4], the actual development of usefully functional systems consisting of large numbers of simple robots will provide ample challenges in both the technical and management dimensions [5]. Not the least of the management challenges is the need to provide a capability to control the system in terms of meaningful mission-oriented system-level parameters. A user (or military "commander") will require an understanding of a system's capabilities, doctrine for employing it, and measures of effectiveness to assess its performance once deployed. It is thus necessary to relate system (ensemble) functionality and performance to the behaviors realized by the individual autonomous elements [5].

Many potential applications for many-robot systems involve the performance of some function which can be characterized as "coverage": the application of the effects of some sensor or effector to some extended physical space. *Blanket* coverage involves a static deployment of robots to effectively detect targets appearing within an extended area; *barrier* coverage minimizes undetected penetration through a barrier of robots; *sweep* coverage moves robots through an area to detect targets [5]. Potential military applications of coverage behaviors that have been identified include intelligent land mine deployment [6], shallow water mine sweeping [7], reconnaissance, sentry duty, communications relay [8,9], maintenance inspection, carrier deck FOD disposal [5], and ship hull cleaning [1].

While the coverage paradigm certainly includes effector-based applications such as painting and cleaning, the canonical coverage application is that of *search*, in which each robotic element carries a sensor which it uses to detect targets of interest, and, in fact, the definitions of blanket, barrier, and sweep coverage presented above are narrowly framed in the language of a search problem. A well established theoretical treatment of search problems has been developed within the operations research community, with roots in the problem of hunting German submarines during World War II [10,11,12]. Grounded in Bayesian decision theory, the field of search theory has, over the past half century, treated a wide variety of problems associated with search (see, for example, the recent survey by Benkoski et al [13]).

A simple search theoretical framework was used in a previous paper [14] to demonstrate that the design of a cost-effective many-robot search system can depend sensitively on the interplay of sensor cost and performance levels with mission-specific functional and performance requirements. In the current paper we extend these results with a more

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE SEP 1993		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Randomized Search Strategies With Imperfect Sensors				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Navy Command, Control and Ocean Surveillance Center Division 531 RDT&E San Diego, CA 92152-7383				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

thorough grounding in optimal search theory. We consider additional measures of effectiveness for area search systems to provide a basis for a tradeoff of coordinated versus random search strategies, and we explore how to deliberately achieve effectively randomized search strategies that provide uniform search coverage over a specified area.

2. ABSTRACTIONS OF TARGET DETECTION SENSORS

We consider the following simple search problem: one or more searchers, each carrying a detection sensor of some sort, move through a predefined search area attempting to detect one or more stationary target objects. The targets are a priori equally likely to be at any point within the search area. Whether any given searcher actually detects any given target at any instant depends on the relative geometry of searcher and target at that instant, and the physical characteristics of the detection sensor(s), the environment, and the target itself.

2.1 Lateral Range Curve

Since, as a rule, detection is more likely the closer the searcher is to the target, it is possible to calculate a cumulative probability of detection over a complete encounter by integrating the instantaneous probability of detection as the searcher approaches a stationary target from far away (where the instantaneous detection probability is very small), passes near the target, and again retreats, as depicted in Figure 1a. Plotting this integrated encounter probability versus the searcher's distance of closest approach to the target (lateral range) yields the *lateral range curve*. In general, a lateral range curve need not be symmetric (a searching ship may have a more sensitive sensor on its port side than on starboard), but in general the values should fall toward zero with increasing lateral range. Figure 1b shows a plausible fictitious lateral range curve.




Figure 1. Lateral range curve and sensor approximations. (a) Derivation of lateral range curve. (b) Plausible lateral range curve. (c) Definite range law approximation, using one parameter. (d) Imperfect sensor approximation, using two parameters.

2.2 Definite Range Law, or "Cookie Cutter"

It is common to approximate a detection sensor's lateral range curve in order to facilitate analytical modeling of the search process. The simplest approximation is that of a *definite range law*: targets which come within a certain distance of the searcher are always detected, and targets which do not come that close are never detected (hence the

searcher cuts a clean swath like a "cookie cutter"). This results in the lateral range curve of Figure 1c, with the critical range value chosen so that the area under the simplified curve is the same as for the "real" sensor. This model thus yields the correct number of targets detected by a single searcher making a single pass through an area. Also, the sensor is characterized by only a single parameter: its maximum detection range.

2.3 Imperfect sensor

A somewhat more sophisticated approximation to a lateral range law results if we are willing to add a second parameter and specify some mean probability of detection less than one over a specified range, resulting in the lateral range curve of Figure 1d. As was the case with the "cookie cutter", this model gives the correct number of targets detected by a single searcher making a single pass through an area. The cookie cutter, however, would give unreasonably optimistic predictions when multiple coordinated passes are made through an area. Since many-robot systems are likely to use low quality sensors, the *imperfect sensor* model is designed to model this situation more reasonably.

3. THE SEARCH PROCESS

Let us now be more specific in describing our search system and application. Our system consists of some number N of identical robotic elements, each of which can move about while carrying an imperfect sensor (in the sense of section 2.3) of nominal range r and detection probability p : any target which lies within a distance r of an element's track is detected with probability p as the element passes by. Targets farther than r are never detected, false alarm detections do not occur, and what action the element might take with detected targets is of no concern. Each element is capable of traveling a total distance d during the mission; this limitation may be due to limited energy storage, or to operational constraints on the duration of the mission, coupled with the maximum speed which can be achieved while operating the sensor payload. Let S be the average number of times that each point in the search area (whose area is A) is sensed; we call S the "sweep fraction" of the search, and calculate it as

$$S = 2r d N / A \quad (1)$$

3.1 Coordinated and Random Search Processes

If we let D represent the probability that any given target is detected, or, equivalently, the expected fraction of targets detected, then our first step is to calculate D as a function of S for two candidate search strategies: in the first case, the elements execute a perfectly coordinated search pattern -- we may have only one element following a "lawn-mower" pattern, perhaps by employing a very accurate navigational system, or we may have a number of elements moving in a tightly coordinated formation; in the second case, we employ less expensive elements, capable only of staying within the designated search area, but otherwise wandering completely randomly.

We can also express these two archetypal search strategies in abstract and discretized form to provide the basis for further probabilistic analysis:

We are given an urn containing a very large number of otherwise identical marbles, a small percentage of which bear an invisible mark. We have a machine that can tell us if a marble presented to it is one that bears a mark. The machine is not perfect: it detects a marked marble only with probability p , but it never indicates a false positive. The task is to separate out the marked marbles by running the marbles through the testing machine. The two candidate strategies are:

Coordinated search: we pull marbles at random from the urn and test them on the machine. After testing, we place the unmarked marbles into a second urn. When all the marbles have been tested, we pour these marbles back into the first urn and repeat the process.

Random search: we pull marbles at random from the urn and test them on the machine. After testing, we return each unmarked marble to the urn and mix it thoroughly.

What can we say about these two abstract strategies? First, the coordinated search is clearly superior, since we never test a marble for the n th time until all marbles have been tested $n - 1$ times (and the more times a marble has been

tested, the less likely that it is an undetected marked marble). If $p = 1$, we will find all the marked marbles by making just one pass (testing each marble just once). On the other hand, with a random search we can never guarantee that we will find all the marked marbles, no matter how many passes we make. The second thing to note is that in order to perform the coordinated search we have to have a second urn; in real life as well, implementing a coordinated search will generally entail additional costs.

In fact, a coordinated search is the best strategy you can employ, and most real world search tactics are designed to achieve a coordinated search. Unfortunately, real world constraints (such as navigation inaccuracies) mean that the results actually achieved usually more closely reflect those of a random search. On the other hand, a random search is not the *worst* strategy that can be employed, and it is not readily obvious how to realize a *deliberately* random search of a two dimensional or three dimensional space.

In the marble world abstraction, S measures the average number of times each marble has been tested: the total number of tests we have made divided by the number of marbles in the urn. For the coordinated case, we calculate

$$D_c = 1 - (1 - p)^{S_c} \quad (2)$$

In fact, this equation is true only for integer S , with straight line interpolation between integers, but we won't worry about this approximation for now -- it provides an upper bound on the performance of the coordinated strategy, and the approximation gets very good for large S . On the other hand, we find for the random case (in the limit of infinitely many marbles)

$$D_r = 1 - e^{-p S_r} \quad (3)$$

Unlike equation (2), equation (3) is accurate for non-integer S . Figure 2 shows the behaviors of equations (2) and (3) when $p = 0.8$.

We now consider the question of how much "better" is the completely coordinated search (first case) than the completely random search (second case).

3.2 Measures of Search Strategy Effectiveness

Different search applications require different measures of systems effectiveness. The task may involve a single target (a sunken submarine, for example), or may involve many targets. And, in the case of many targets, the desire to maximize target detections per amount of search effort (as in prospecting for manganese nodules) is not the same thing as the desire to minimize the number of targets missed in a swept area (as in minesweeping). Finally, a search problem may be embedded in a larger optimization context, as in "what is the best strategy for two adults to use in searching for a stationary lost child so as to minimize the expected time until all three are reunited?" [15]

In [14], we considered a minesweeping application in which the goal was to minimize the cost of detecting a specified fraction D of the targets. By equating D_c in equation (2) with D_r in equation (3), we calculated the "search gain" G , or factor reduction in required search effort afforded by a coordinated search as compared to a random search:

$$G_{\text{many target}} = S_r / S_c = -\ln(1 - p) / p \quad (4)$$

This value, which is independent of the specified required detect fraction D , is plotted in Figure 3. We used this result in [14] as the basis for a "toy" system design tradeoff exercise, configuring an "optimal" system from candidate component subsystems of various capabilities and costs. The design space consisted of 12 possible systems, allowing any of three possible sensors, a completely random or a completely coordinated search strategy, and an optional battery upgrade to double vehicle range, with each choice having an associated specified cost. As the quality of the sensor (its raw target detection probability p) improved, the most cost effective design shifted from employing a large number of the least expensive and least capable elements to a much smaller number of more expensive and more capable elements.




Figure 2. Detect Fraction (D) as a function of Sweep Fraction (S) for sensor probability of detection (p) equal to 0.8. (a) Completely coordinated search strategy, equation (2). (b) Completely random search strategy, equation (3).

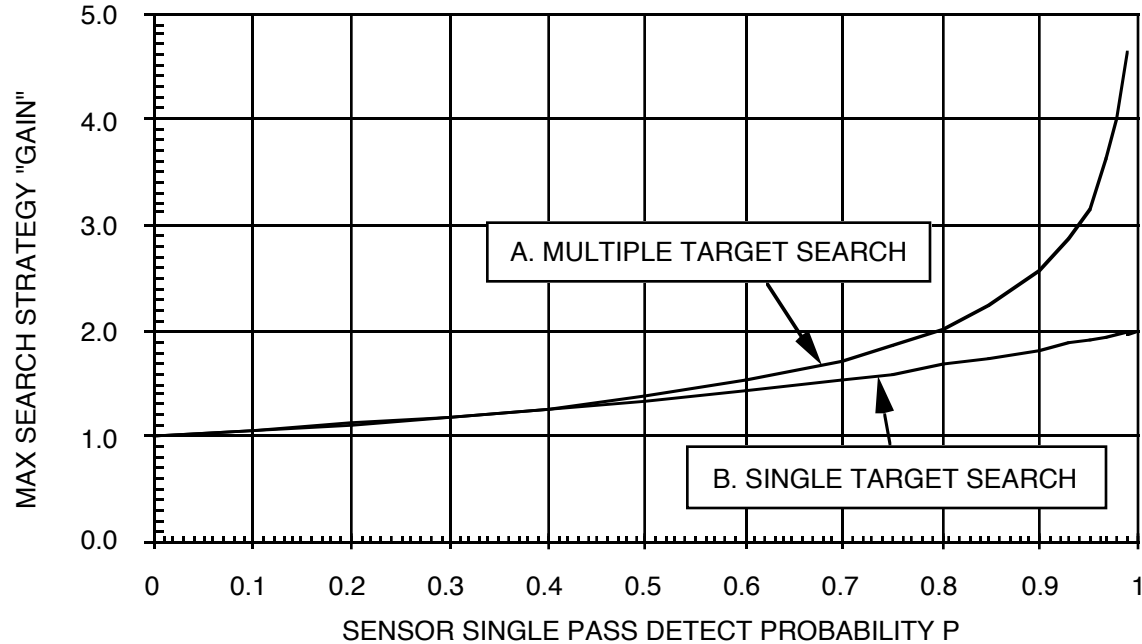


Figure 3. Search strategy gain (G) of a completely coordinated search strategy over a completely random strategy, as a function of the sensor target detection probability (p), for (a) finding a specified fraction of multiple targets, equation (4), and (b) finding a single target, equation (9).

Here we present the analogous analysis for the case of a search for a single target. Now the goal is to minimize the expected search effort required to find the target. Again letting S be the sweep fraction or search effort variable, and letting $\text{Prob}(S)$ be the probability that the single target is discovered after the expenditure of effort S , we can write:

$$\langle S \rangle = \int_0^\infty S \text{Prob}(S) dS \quad (5)$$

For a random search, we have

$$\langle S_r \rangle = \int_0^\infty S e^{-pS} dS = 1/p \quad (6)$$

which is a well known result. Also well known is the result for a coordinated search when $p = 1$

$$\langle S \rangle = \int_0^1 S dS = 1/2 \quad (7)$$

Since the value of $\langle S \rangle$ for a random search with $p = 1$ is just 1, we see that, for a cookie cutter sensor, a coordinated search for a single target is twice as good as a random search. Less well known, however, is the formula for a coordinated search when p is less than 1. With a suitable substitution of variables, the formula is

$$\langle S_c \rangle = \sum_{n=0}^{\infty} \int_0^1 (n+x) p (1-p)^n dx = 1/p - 1/2 \quad (8)$$

The "search gain" due to a fully coordinated search for a single target, compared to a random search, is thus

$$G_{\text{single target}} = \langle S_r \rangle / \langle S_c \rangle = (1/p) / (1/p - 1/2) = 2 / (2 - p) \quad (9)$$

This value is also plotted in Figure 3. With an imperfect sensor, then, we see that the advantage of a coordinated search for a single target is close to that for the multiple target search when p is low, but is limited to a maximum value of 2 when $p = 1$. Cost / benefit tradeoffs can be done analogous to those presented in [14] for the multiple target search, and the results will be more favorable to the random search. The point is that, since the G value is different for the two applications, the optimal choice of sensor, search strategy, and other aspects of a system may well be different as well. It is entirely possible to devise two search systems A and B such that A will find any specified percentage of a field of mines with less cost than B, while B will, on average, find a single mine with less cost than A.

4. RANDOMIZED SEARCH STRATEGIES

In the previous section we considered random search from the perspective of the performance predicted by a particular analytical model; in this section we consider the problem of how to generate a randomized search path which provides uniform coverage over a designated search area.

By "randomized" we mean a path which is not predictable, so that a mobile target which somehow gains partial or even complete knowledge of a searcher's prior path and current position can not predict its future path and thereby evade detection. The strategy must clearly not attempt to be more efficient than the random search model, since in this case the target could evade capture by favoring locations that the searcher had already visited. In general, the searcher's path must not follow any predictably discernible pattern that can permit the target to evade capture. More critically, the search path must not be determined by fixed reaction to environmental features, since this may permit the target to spoof the searcher into changing his path (see [16] for a discussion of security issues relevant to robotic systems). In addition, navigational inaccuracies due to sensor bias may lead to systematic gaps in coverage even if the target is not sophisticated.

Use of a *diffusion* strategy has been analyzed by Eagle at the Naval Postgraduate School [17] and, more recently, has been simulated in the context of a many-robot system for neutralizing shallow water mines [7]. A strategy of following straight line paths between points chosen at random within the search area was proposed by Henze [18], but simulations by McNish [19] showed that this does not lead to a uniform distribution of search effort, instead providing excessive coverage of the central area at the expense of the periphery. In order to increase the coverage near the boundaries, McNish considered a number of other algorithms that generate paths consisting of sequences of chords within a circular search area. The use of a search path consisting of chords is appealing for several reasons. First, the searcher travels as far as possible between changes of direction and is guaranteed not to visit any point twice during the transit of a chord. Second, a chord-based strategy is completely specified by the algorithm that determines the direction of the next chord each time the searcher arrives at the boundary of the search area. The specific reflection algorithms simulated by McNish were *specular* reflection (angle of reflection equal to angle of incidence, like light from a mirror), *uniform* reflection (angle of reflection distributed uniformly), and *diffuse* reflection (like light reflecting from a matte surface). McNish found that only diffuse reflection reliably provided uniform search coverage.

4.1 Diffuse Reflection in a Circular Search Area

The diffuse reflection algorithm specifies a random distribution of reflection angle θ with probability defined as:

$$\text{Prob}(\theta) = 1/2 \sin(\theta) d\theta \quad (10)$$

where θ measures the angle of the chord from the tangent to the boundary at the reflection point. Lalley and Robbins [20] mathematically proved that, in the limit as the length of the search path goes to infinity (corresponding to the sensor range going to 0), the search coverage is uniformly distributed over the circular disk: "the long-run search effort devoted to any region of D is proportional to its area, and consequently, the 'infinitesimal search effort' is the same at every point of D . It is intuitively plausible that a good search plan should have this property, for if not, a target could gain an advantage by hiding at a point scheduled for a low level of search effort" [20].

4.2 Generalization to Convex Search Areas

There are several reasons why it is highly desirable to extend the diffuse reflection strategy to generalized convex search areas. For example, it is impossible to tile a space with circles, so a strategy for a rectangular area would be nice if it is desired to divide a large search area up into smaller subareas (to allocate the search among many searchers, or perhaps to apply additional search effort in regions where targets are a priori more likely to be found).

Lalley and Robbins are apparently of two minds on the question of whether diffuse reflection applies unchanged to the general convex case. In [20] they categorically state: "the search plan studied here is suitable for only circular domains; there is no immediate generalization to other shapes", while in [21] they prove that the diffuse reflection strategy is asymptotically minimax as the sensor range goes to 0 in the two player game known as "Princess and Monster", in which the Princess (the target) is allowed to move about in the search space, but has no information about the Monster's position [22]

We present here a straightforward generalization of diffuse reflection that applies to a search for stationary targets when the searcher knows the boundaries of the search area and his position within it.

Referring to Figure 4a, we choose any point P on the boundary of the convex search area D . Let θ measure the angle from the tangent to the boundary of D at P to any chord, and let $K(\theta)$ measure the length of the chord from P across D at angle θ to the tangent. Now consider a wedge shaped region of the convex search space D formed by two chords PQ and PR drawn from P at angles θ_1 and θ_2 respectively, and the boundary segment RQ opposite P . We can now write:

$$\text{Area}(A) = \int_{\theta_1}^{\theta_2} 1/2 K(\theta)^2 d\theta \quad (11)$$

$$\text{Effort}(\theta) = K(\theta) \text{ width} \quad (12)$$

$$\langle \text{Effort}(A) \rangle = \int_{\theta_1}^{\theta_2} \text{Prob}(\theta) \text{Effort}(\theta) d\theta \quad (13)$$

The solution for Prob (q) follows immediately from Lalley and Robbins' observation quoted above; that, in a good search plan, the "search effort devoted to any region of D is proportional to its area". Substituting equation (12) in equation (13) we see that if $\langle \text{Effort}(A) \rangle$ is to be proportional to Area (A) for any θ_1 and θ_2 then

$$\text{Prob}(\theta) = \text{constant } K(\theta) \quad (14)$$




Figure 4. $K(\theta)$ generalization of diffuse reflection strategy. (a) Derivation of $K(\theta)$ formula, equation (14). (b) Demonstration of equivalence of $K(\theta)$ with diffuse reflection for case of a circular disk.

This algorithm is "optimal" in that it provides a uniform expectation of search effort over the entire area on a per-chord basis. In order to implement this algorithm, however, the searcher must be able to determine his position within the search area, and the distance to the opposite boundary in all directions. This will be trivial in some applications, but impossible in others. As Figure 4b demonstrates, application of the $K(\theta)$ strategy to the special case of a circular disk does indeed yield the diffuse reflection law, and, since all points on the circumference are equivalent, the searcher need only be capable of determining the tangent to the boundary at his position, and does not need to know his location on the circumference. In the absence of the information required to implement the $K(\theta)$ algorithm, a searcher within a convex region should use the diffuse reflection algorithm.

It is not possible to extend the strategy to generalized nonconvex areas, since there will be some points in the search space from which not all other points are directly reachable by a chord trajectory. However, the strategy can be trivially extended to a nonconvex space that is made up of two joined mirror-symmetric convex spaces (such as a Valentine heart).

5. CONCLUSION

Search systems consisting of many inexpensive robots will tend to use randomized rather than coordinated search strategies for two reasons: (1) the increase in effectiveness provided by a coordinated (as opposed to random) search strategy decreases as the capability of the search sensor (probability of target detection) decreases; and (2) the cost of implementing the navigation capabilities necessary to support a coordinated search strategy may be prohibitive, relative to the cost of a less capable search element. However, careful analysis is required to select and implement an appropriate strategy that efficiently provides uniform area coverage. Developers of robotic search systems must do more than simulate one or two candidate search strategies; fifty years of developments in the field of search theory can provide them with analytical tools to help understand tradeoffs relative to overall system search performance. Meanwhile, the advent of search systems consisting of inexpensive robots carrying minimum capability sensors provides search theorists with the opportunities of a previously unexplored problem domain.

6. ACKNOWLEDGEMENTS

This work is supported by the Intelligent Systems Program of the Office of Naval Research, Computer Science Division. The author wishes to acknowledge stimulating interactions with Alan Washburn and Jim Eagle of the Naval Postgraduate School, Jill Crisman of Northeastern University, Jerome Kirsch of Grumman Electronics Systems, Dick Blidberg of the UNH Marine Systems Engineering Laboratory, and Teresa McMullen of ONR.

7. REFERENCES

1. Flynn, A.M. "Gnat Robots (And How They Will Change Robotics)", **Proceedings of the IEEE Micro Robots and Teleoperators Workshop**, Hyannis MA, 9-11 November 1987. Also appeared in **AI Expert**, December 1987, p 34 et seq.
2. Freedman, D.H. "Invasion of the Insect Robots", **Discover**, March 1991, pp 42-50.
3. Flam, F. "Swarms of Mini-Robots Set to Take on Mars Terrain", **Science**, 18 September 1992, p 1621.
3. "Ants under Study as Model for Small, Low Cost Robots", **Richmond Times-Dispatch**, 21 June 1992.
5. Gage, D.W. "Command Control for Many-Robot Systems", **AUVS-92, the Nineteenth Annual AUVS Technical Symposium**, Huntsville AL, 22-24 June 1992. Reprinted in **Unmanned Systems Magazine**, Fall 1992, Volume 10, Number 4, pp 28-34.
6. Macedonia, R.M., and B. Sert. "Low Cost Mobility for Robotic Weapons", **AUVS-88, the Fifteenth Annual AUVS Technical Symposium**, San Diego CA, 6-8 June 1988.
7. Mangolds, A. "Lemmings - a Swarming Approach to Shallow Water Mine Field Clearance", **AUVS-93, the Twentieth Annual AUVS Technical Symposium**, Washington DC, 29-30 June 1993, pp xx.
8. Megatek Corporation, "CRICKET: A Novel LPI Communication Concept", Report R2017-003-IF-2, 30 September 1978.
9. Kirsch, J. "Obtaining Group Synergy Using Swarms of Small Expendable Tactical Robots", **AUVS-93, the Twentieth Annual AUVS Technical Symposium**, Washington DC, 29-30 June 1993, pp xx.
10. Koopman, B.O. **Search and Screening**. Pergamon Press, New York, 1980.
11. Stone, L.D. **Theory of Optimal Search**, Academic Press, New York, 1975.
12. Washburn, A.R. **Search and Detection**, Operations Research Society of America, 1981.

13. Benkoski, S.J., Monticino, M.G., and Weisinger, J.R. "A Survey of the Search Theory Literature", **Naval Research Logistics**, Volume 38, August 1991, pp 469-494.
14. Gage, D.W. "Sensor Abstractions to Support Many-Robot Systems", **Proceedings of SPIE Mobile Robots VII**, Boston MA, 18-20 November 1992, Volume 1831, pp 235-246.
15. Thomas, L.C., "Finding Your Kids When They Are Lost", **Journal of the Operations Research Society**, Volume 43, Number 6, June 1992, pp 637-639.
16. Gage, D.W. "Security Considerations for Autonomous Robots", **Proceedings of Symposium on Security and Privacy**, Oakland CA, 22-24 April 1985, pp 224-228. Reprinted in **Computer Security Journal**, vol 6, 1990, pp 95-99.
17. Eagle, J.N. "Estimating the Probability of a Diffusing Target Encountering a Stationary Sensor", **Naval Research Logistics**, Volume 34, 1987, pp 43-51.
18. Henze, J. "Random Search Formulas and the Distribution of Detection Times", unpublished study for COMSUBPAC, August 1986.
19. McNish, M.J. "Effects of Uniform Target Density on Random Search", Master's Thesis, Naval Postgraduate School, September 1987.
20. Lalley, S.P., and Robbins, H.E. "Uniformly Ergodic Search in a Disk". **Lecture notes in pure and applied mathematics, v. 112, Search Theory**. Marcel Dekker, Inc., New York, 1989, pp 131-151.
21. Lalley, S.P., and Robbins, H.E. "Stochastic Search in a convex region", **Probability Theory and related Fields**, vol 77, 1988, pp 99-116.
22. Fitzgerald, C. "The Princess and Monster Differential Game", **SIAM Journal on Control and Optimization**, Volume 17, 1979, pp 700-712.